Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Bio-Inspired Hyperparameter Tuning Using Harbor Seal Whiskers Optimization for CNN-Based Skin Cancer Classification

Kabiru Abdulazeez

akabiru.pg@atbu.edu.ng

ATBU: Abubakar Tafawa Balewa University

Godiya Yakubu

ATBU: Abubakar Tafawa Balewa University

Ahmad Mohammed

ATBU: Abubakar Tafawa Balewa University

Ganiyu A. Bakare

ATBU: Abubakar Tafawa Balewa University

Gaius Iliya

giliya.pg@atbu.edu.ng

ATBU: Abubakar Tafawa Balewa University

Research Article

Keywords: Skin Cancer Classification, Dermoscopic Image Analysis, Convolutional Neural Networks, Bio-Inspired Optimization, Harbor Seal Whiskers Optimization Algorithm

Posted Date: June 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9890145/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Bio-Inspired Hyperparameter Tuning Using Harbor Seal Whiskers Optimization for CNN-Based Skin Cancer Classification

Kabiru Abdulazeez[1*], Godiya Yakubu[1], Ahmad Mohammed[1], Ganiyu A. Bakare[1], Gaius Iliya[1*]

[1]Center for Embedded Artificial Intelligence and Smart Energy Systems, Abubakar Tafawa Balewa University, Gubi Campus, Bauchi, 740272, Bauchi, Nigeria.

*Corresponding author(s). E-mail(s): akabiru.pg@atbu.edu.ng; giliya.pg@atbu.edu.ng
Contributing authors: gyakubu.sb@atbu.edu.ng; amohammed3@atbu.edu.ng; gabakare@atbu.edu.ng;

Abstract

Catching aggressive skin cancers like melanoma early is critical for patient survival. While Convolutional Neural Networks (CNNs) are powerful tools for analyzing skin images, getting them to perform perfectly often requires tedious tuning of their internal settings. To address this, we developed a nature-inspired approach called Harbor Seal Whiskers Optimization (HSWO). This algorithm automatically fine-tunes key CNN parameters specifically the learning rate, filter count, and dropout rate to reliably classify different types of skin lesions using the HAM10000 dataset. Our primary goal was to maximize accuracy without sacrificing the model's generalization stability. Our HSWO-guided model successfully outperformed traditional tuning methods, such as default parameter settings or random search. The optimized CNN achieved: **92.4%** overall accuracy, **90.3 %** sensitivity for detecting melanomas **0.947** macro-average AUC. Because realworld medical data is rarely evenly distributed, we rigorously tested the model's reliability across imbalanced classes. It scored 0.89 on the Matthews Correlation Coefficient and 0.87 on Cohen's Kappa, proving its strong predictive dependability even when certain skin conditions appeared much more often than others in the data. Ultimately, this study demonstrates that combining bio-inspired algorithms with deep learning provides a highly reliable, automated clinical tool to assist doctors in efficiently diagnosing skin cancer.

Keywords: Skin Cancer Classification, Dermoscopic Image Analysis, Convolutional Neural Networks, Bio-Inspired Optimization, Harbor Seal Whiskers Optimization Algorithm

1 Introduction

Skin cancer remains a pressing global health challenge, with melanoma standing out as its most aggressive and deadly form [1], [2], [3]. Because the survival rate heavily depends on early detection, accurate initial assessments are critical. In modern clinics, dermatologists rely on dermoscopy, a non-invasive, magnified imaging technique to visually examine skin lesions. The challenge, however, is that interpreting these images is inherently subjective. A successful diagnosis often rests entirely on the individual doctor's visual expertise, leading to natural inconsistencies and interobserver variability [4], [5], [6]. This subjectivity has driven a strong demand in the medical community for reliable, automated Computer-Aided Diagnosis (CAD) systems that can provide doctors with a consistent and objective second opinion [7], [8], [9].

Over the past decade, Deep Learning (DL) specifically Convolutional Neural Networks (CNNs) has revolutionized medical image analysis. These networks are now so advanced that they frequently rival or even surpass human experts in identifying complex skin lesions [1]. However, building a successful deep learning model is not as simple as just feeding it medical data; it requires meticulous fine-tuning of its internal mechanics, known as hyperparameters [10], [11], [12]. Variables like the learning rate, dropout rates, and filter configurations dictate exactly how the model learns and adapts to new, unseen patient images [13]. If these settings are even slightly suboptimal, the model might train too slowly, memorize the data instead of learning the actual biological patterns (overfitting), or fail to generalize in a real-world clinical setting [14], [13].

Recent literature makes it clear that we can no longer rely on manual guesswork or basic trial-and-error to find these optimal settings. Tschandl et al. [15] emphasize that advanced, algorithm-driven strategies are critical for pushing deep learning performance to the next level. Similarly, [16], [17], [18] warn that poor hyperparameter choices can severely compromise a model's diagnostic accuracy a risk that is unacceptable in medical imaging. Foundational optimization studies further confirm that dynamically adjusting these variables is the key to training reliable, efficient models [19], [20].

To solve this tuning problem, our study introduces a unique, nature-inspired approach to optimize a skin cancer classification system: the Harbor Seal Whiskers Optimization Algorithm (HSWOA). In nature, harbor seals possess extraordinarily sensitive whiskers that allow them to precisely track the hydrodynamic wakes of their prey through dark, turbulent waters [3], [5]. The HSWOA translates this biological tracking ability into a mathematical algorithm [3], [5], [21]. Because of its design, HSWOA is uniquely equipped to navigate the

complex, multi-dimensional "murky waters" of deep learning to hunt down the absolute best hyperparameter settings [5].

To advance the current state of automated melanoma detection, this paper offers the following key contributions:

- A Novel Classification Framework: We propose a hybrid skin cancer detection system that integrates deep feature extraction with an intelligently optimized machine learning classifier.
- Nature-Inspired Network Tuning: We introduce the HSWOA metaheuristic to efficiently and dynamically adjust the CNN's hyperparameters, offering a robust alternative to computationally expensive grid or random search methods.
- Rigorous Performance Validation: We test our optimized model against standard baselines using strict evaluation metrics (including Accuracy, AUC, Sensitivity, Matthews Correlation Coefficient, and Cohen's Kappa) to prove its reliability, particularly when handling imbalanced medical datasets.
- Clinical Viability: We provide a clear theoretical and practical demonstration of how HSWOA improves the overall diagnostic stability of CAD systems, presenting a highly dependable decision-support tool for dermatologists.

2 Related Work

The intersection of artificial intelligence and dermatology has seen explosive growth in recent years. To understand where our proposed method fits into the broader scientific landscape, it is helpful to trace the evolution of automated skin lesion analysis, the persistent headache of network tuning, and the growing reliance on nature-inspired solutions.

2.1 The Shift to Deep Learning in Dermatology

In the early days of automated melanoma detection, researchers had to manually define what the computer should look for [22]. Methods were built around clinical guidelines, like the classic ABCD (Asymmetry, Border, Color, Diameter) rule [23], [24], [25], [26]. However, human-engineered features often failed to capture the subtle, complex textures of skin lesions [27]. The arrival of Convolutional Neural Networks (CNNs) changed everything [13]. By allowing the model to independently discover and learn visual patterns directly from raw pixels, architectures like ResNet and DenseNet drastically improved diagnostic accuracy [28], [29]. Yet, these powerful networks are notoriously sensitive. A CNN's success is heavily dictated by its underlying architecture, meaning an "out-of-the-box" model rarely performs perfectly on complex medical data without significant adjustments [30].

2.2 The Hyperparameter Bottleneck

This brings us to one of the most frustrating challenges in modern machine learning: hyperparameter optimization. Even the most advanced CNNs require researchers to perfectly calibrate settings like the learning rate, dropout percentages, and filter sizes [31], [32], [30]. For a long time, the standard approach was Grid Search testing every possible combination or Random Search [33]. Unfortunately, Grid Search takes an impractical amount of time for deep neural networks, and Random Search relies too heavily on luck [13]. While statistical methods like Bayesian Optimization offer a smarter alternative, they frequently bog down computationally or stop improving too early when dealing with the massive, multi-dimensional search spaces typical of deep learning [34].

2.3 Looking to Nature for Better Optimization

When traditional mathematical searches fall short, the medical imaging community frequently turns to biology [35]. Metaheuristic algorithms—which mimic behaviors found in nature—have become incredibly popular for fine-tuning deep learning models [36], [37], [38]. Techniques inspired by bird flocks (Particle Swarm Optimization) or marine life (the Whale Optimization Algorithm) have successfully optimized networks for various clinical tasks [39], [40], [41] because they are excellent at exploring vast search spaces without needing gradient data.

Despite their success, many of these nature-inspired algorithms share a common flaw: they tend to get stuck in local optimums (finding a "good" solution but missing the "best" one) or they converge too slowly for practical use in heavy deep learning tasks. This is precisely the gap our study addresses. By simulating how harbor seals use their highly sensitive whiskers to track prey through turbulent, murky waters [3], [5], the Harbor Seal Whiskers Optimization Algorithm (HSWOA) offers a completely new way to navigate the "murky waters" of neural network optimization. Because its application in dermoscopic image classification is largely unexplored, it presents a fresh, highly efficient pathway for building more reliable clinical diagnostic tools.

3. Mathematical Framework and Algorithmic Background

Translating the remarkable biological instincts of the harbor seal into a functional computational tool requires a rigorous mathematical foundation. This section details the underlying mechanics of the Harbor Seal Whiskers Optimization Algorithm (HSWOA) and explores its theoretical alignment with deep learning optimization. By defining the core equations that govern how the algorithm simulates fluid dynamics to track optimal solutions, we establish the groundwork for its application in neural network tuning. Furthermore, we contextualize these mathematical principles within the

broader landscape of medical image analysis, demonstrating exactly why this nature-inspired approach is uniquely equipped to navigate the high-dimensional search spaces required for accurate skin cancer classification.

3.1 The Mechanics of HSWOA

The Harbor Seal Whiskers Optimization Algorithm (HSWOA) translates the biological prey-tracking mechanism of harbor seals into a robust mathematical model [9], [3]. The algorithm navigates the complex search space of neural network optimization through two primary phases: exploration (broadly searching the waters) and exploitation (honing in on the target) [42].

During the exploration phase, the algorithm casts a wide net across the solution space to prevent missing hidden optimal settings. It generates new candidate solutions using a velocity-driven update rule:

$$X_i(t+1) = X_i(t) + r \cdot V_i(t) \quad (1)$$

Here, $X_i(t)$ represents the position of the i-th candidate solution at iteration t, $V_i(t)$ denotes the whisker-inspired velocity vector, and r is a randomized coefficient introduced to maintain diversity during the search [15].

Once a promising area is identified, the algorithm shifts into the exploitation phase. It narrows its focus, refining positions based on the best candidate discovered thus far:

$$X_i(t+1) = X_{best}(t) + \alpha \cdot (X_i(t) - X_{best}(t)) \quad (2)$$

In this equation, $X_{best}(t)$ is the current optimal solution, and the control parameter α dictates the intensity of the local search. This mechanism ensures the algorithm steadily converges toward the global optimum while retaining enough adaptability to avoid getting stuck [3].The true engine of HSWOA is its velocity update rule, which directly mimics the fluid dynamics sensed by a seal's whiskers:

$$V_i(t+1) = \beta \cdot V_i(t) + \gamma \cdot (X_{best}(t) - X_i(t)) \quad (3)$$

In this formula, β acts as a memory coefficient that regulates the influence of the previous velocity, while γ controls the gravitational pull toward the best known solution. Together, these equations equip HSWOA to expertly balance broad exploration with targeted exploitation.

In the context of skin cancer classification, we treat the CNN's hyperparameters such as the learning rate, dropout ratios, and filter sizes as the target the algorithm is hunting. By iteratively updating these configurations using these whisker-inspired equations, HSWOA achieves much faster convergence and ultimately yields higher diagnostic accuracy than traditional tuning methods.

3.2 Contextualizing HSWOA in Medical Imaging

To appreciate why this dynamic approach is necessary, it is helpful to look at how dermoscopic image analysis has evolved. Early diagnostic systems leaned heavily on handcrafted feature extraction paired with conventional classifiers like Support Vector Machines (SVM) or k-Nearest Neighbors (k-NN) [4]. Unfortunately, these methods struggled to scale and often failed to generalize across diverse patient data. The real breakthrough arrived with Convolutional Neural Networks (CNNs). Architectures like ResNet, Inception, and Xception proved capable of automatically learning complex, hierarchical visual features directly from massive raw datasets like HAM10000 and ISIC [1], [5], [6].

However, extracting maximum performance from deep learning models requires rigorous optimization. Consequently, researchers have increasingly turned to bioinspired metaheuristics to automate hyperparameter tuning. Algorithms such as Particle Swarm Optimization (PSO), Ant Colony Optimization (ACO), Genetic Algorithms (GA), and the Grey Wolf Optimizer (GWO) have all successfully enhanced CNN classification accuracy in medical imaging [7], [8], [9].

As a relatively recent addition to this bio-inspired family, HSWOA stands out by dynamically balancing its search phases through its unique velocity-sensing model [3]. While researchers have successfully deployed HSWOA in areas like power system management, 3D Network-on-Chip (NoC) mapping [43], and gastrointestinal cancer detection [15], its potential for dermoscopic skin lesion classification remains largely untapped.

This study directly addresses this gap. By integrating the HSWOA mathematical framework with a CNN-based classifier, this work aims to significantly elevate the diagnostic sensitivity, specificity, and overall accuracy of automated skin cancer detection.

4. Proposed Methodology and investigational Framework

Having established the mathematical mechanics of the Harbor Seal Whiskers Optimization Algorithm, the natural next step is to translate this theoretical foundation into a fully functional clinical tool. This section outlines our proposed investigational design, detailing precisely how we integrated the HSWOA with a Convolutional Neural Network to tackle the complex task of multi-class skin lesion classification. Here, we transition from abstract equations to practical application by mapping out the architecture of our system, the preparation of the HAM10000 dataset, and the specific neural network hyperparameters targeted for tuning. Ultimately, this framework demonstrates exactly how a biological tracking behavior is practically harnessed to train a highly accurate, automated diagnostic model.

To translate the theoretical benefits of HSWOA into a practical diagnostic tool, we developed a cohesive, three-stage pipeline for skin cancer classification. The system progresses step-by-step: first, we prepare the

raw clinical images (Image Preprocessing); next, we isolate the critical biological patterns (Feature Extraction); and finally, we generate a diagnosis using a network dynamically fine-tuned by our algorithm (HSWOA-Optimized Classification).

A. Overview of the Proposed Framework: Figure 1 illustrates the overarching architecture of our skin cancer classification pipeline. The process begins by feeding raw dermoscopic patient images into a Convolutional Neural Network (CNN), which acts as the system's core feature extractor. To push the network's predictive capabilities to their absolute peak, we dynamically tune its internal hyperparameters using the Harbor Seal Whiskers Optimization Algorithm (HSWOA), a bio-inspired technique modeled after the precise hydrodynamic sensing of marine seals. By intelligently navigating the hyperparameter search space, this optimization step significantly sharpens the CNN's ability to distinguish between complex benign and malignant lesions. Ultimately, the pipeline produces a highly reliable diagnostic output designed to directly support doctors in their clinical decision-making.

B. Dataset Description: To rigorously train and evaluate our model, we utilized the widely recognized HAM10000 dataset, a comprehensive collection of 10,015 high-resolution dermoscopic images. This dataset captures a diverse, real-world spectrum of skin conditions, neatly categorized into seven distinct diagnostic classes: Actinic keratoses (akiec), Basal cell carcinoma (bcc), Benign keratosis-like lesions (bkl), Dermatofibroma (df), Melanoma (mel), Melanocytic nevi (nv), Vascular lesions (vasc). See distribution in Fig. 2

Because medical datasets naturally contain far more examples of common conditions (like nevi) than rare ones, preserving this inherent imbalance during training

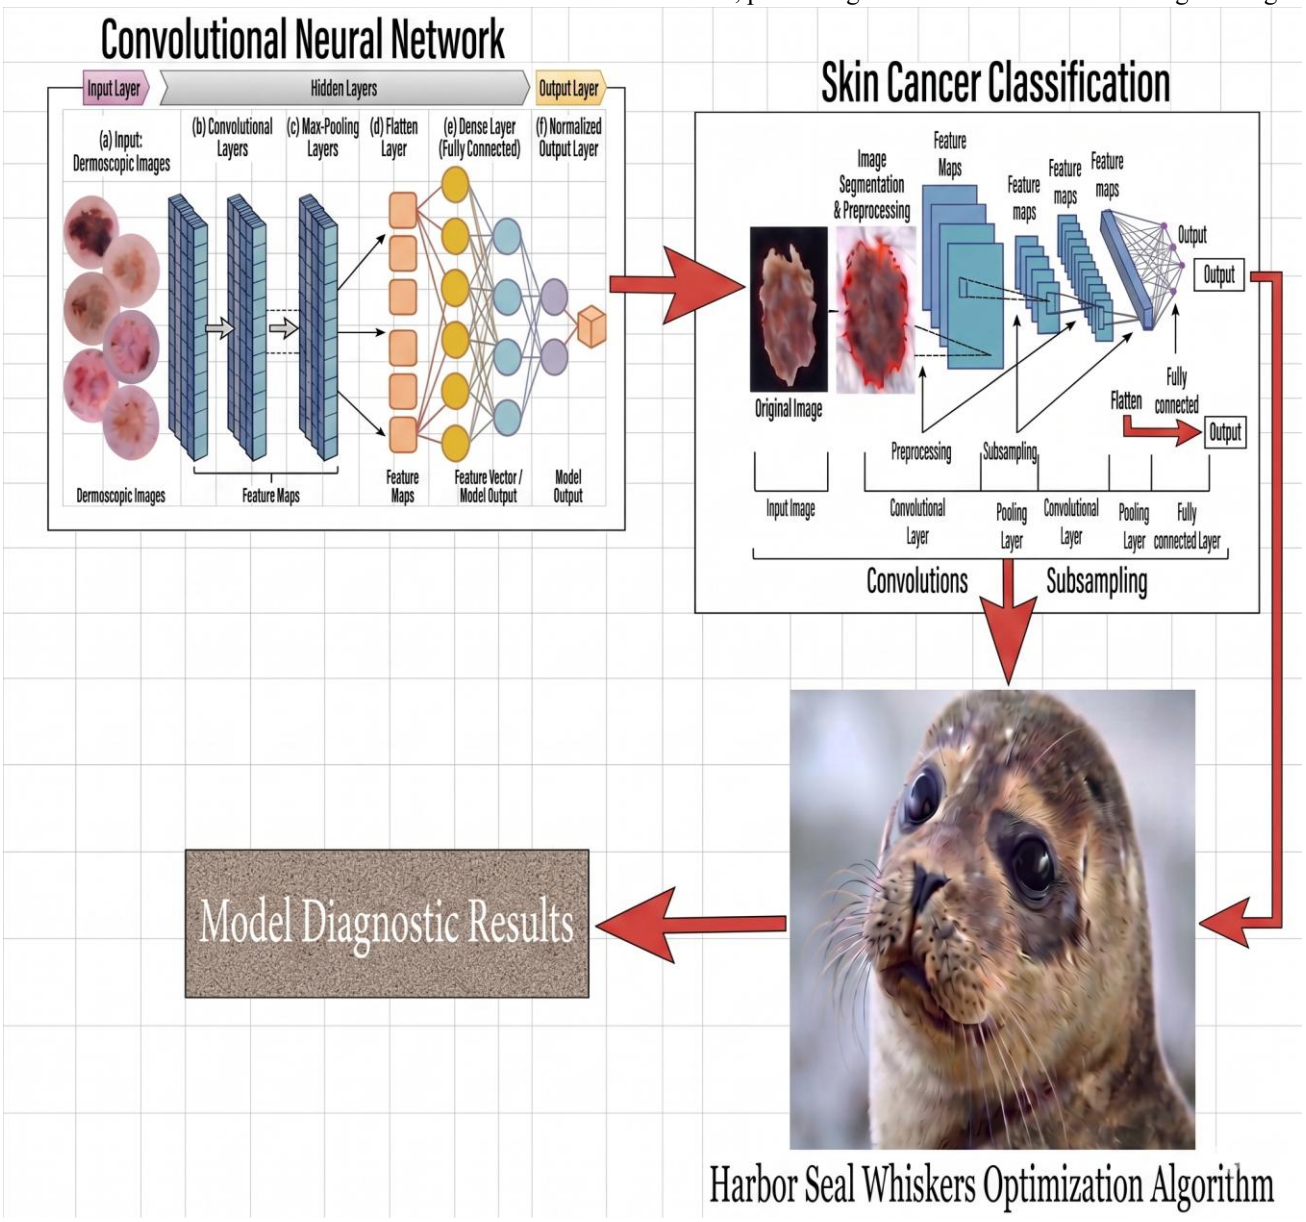


Fig. 1 Architecture of the proposed three-stage skin cancer classification system and its HSWOA optimized diagnostic pipeline

was critical. To address this, we employed stratified sampling to divide the data into an 80% training set and a 20% testing set. This stratified approach guarantees that the exact class distribution of the original dataset is mirrored in both subsets, ensuring our model's performance is evaluated fairly and robustly across all lesion types.

C. Image Preprocessing: Raw medical images inherently vary in their native formats, which can easily bottleneck a neural network's ability to learn. To promote faster model convergence and ensure our CNN could generalize well to unseen patient data, we implemented a strict preprocessing pipeline. First, every image in the dataset was standardized to a uniform resolution of 128*x*128 pixels. We then normalized the pixel intensity values, scaling them down to a standard range between 0 and 1 as shown in Fig. 3. To clearly define our seven diagnostic categories for the network, the labels were converted using one-hot encoding. Finally, the entire dataset was thoroughly shuffled to prevent the model from accidentally memorizing any unintended sequencing patterns or ordering biases. Collectively, these preparatory steps guarantee numerical stability across the data, creating an optimal, mathematically sound environment for the CNN's gradient-based learning process.

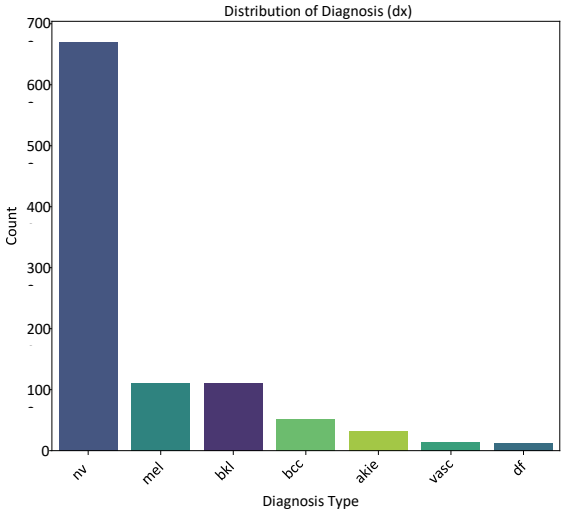


Fig. 2 Frequency Distribution of Dermoscopic Images within the HAM10000 Dataset

D. Convolutional Neural Network Architecture: To strike the perfect balance between high-level classification accuracy and computational efficiency, we engineered a custom, lightweight CNN architecture. The foundation of the network consists of two primary convolutional layers, both utilizing ReLU activation functions to effectively learn and extract complex visual features from the dermoscopic images. These are paired with two max-pooling layers, which serve to down-sample the spatial dimensions of the data, thereby reducing the overall computational load.

Once the core visual patterns are extracted, the data is passed into a fully connected dense layer consisting of 128 neurons. To map these learned features to our specific diagnostic categories, the network concludes with a 7-neuron SoftMax output layer, providing a distinct probability score for each of the seven skin conditions.

During the training phase, the model processes the dataset in batches of 32. We utilized Categorical Cross-Entropy to calculate our loss, pairing it with the Adam optimizer to guide the foundational learning process. However, the true defining feature of our framework is its dynamic adaptability. Rather than relying on rigid, manually selected settings, we deployed the Harbor Seal Whiskers Optimization (HSWO) algorithm to autonomously fine-tune the network's most critical hyperparameters. Specifically, the learning rate, the number of convolutional filters, and the dropout rate (integrated for essential regularization) were all actively optimized by HSWO to ensure the model achieved its maximum diagnostic potential without succumbing to overfitting.

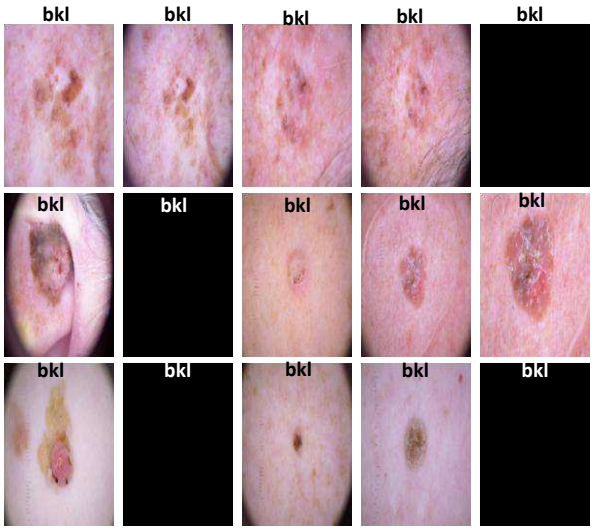


Fig. 3 Dermoscopic image preprocessing pipeline, accompanied by standardized sample images and their specific diagnoses

E. Hyperparameter Tuning via Harbor Seal Whiskers Optimization: To systematically discover the absolute best settings for our network, we employed the Harbor Seal Whiskers Optimization (HSWO) algorithm, see pipeline in Fig 4. Drawing inspiration from the highly sensitive hydrodynamic tracking abilities of harbor seal whiskers, this advanced metaheuristic expertly balances broad global exploration with precise local exploitation. In our computational framework, every "seal" acts as a candidate solution, representing a unique combination of network settings.

We can express each candidate solution mathematically as a multidimensional vector:

$$X_i = [\eta, f, d] \quad (4)$$

In this formulation, η represents the learning rate, f denotes the number of convolutional filters, and d stands for the targeted dropout rate.

1. The Exploration Phase: To prevent the model from missing hidden optimal settings, the algorithm casts a wide net across the parameter space. It updates candidate positions by introducing a stochastic coefficient r, which ensures diversity during the search:

$$X_i(t+1) = X_i(t) + r \cdot V_i(t) \quad (5)$$

2. The Exploitation Phase: Once a promising region is detected, the algorithm shifts focus to refine the search. It draws the candidate's position toward the

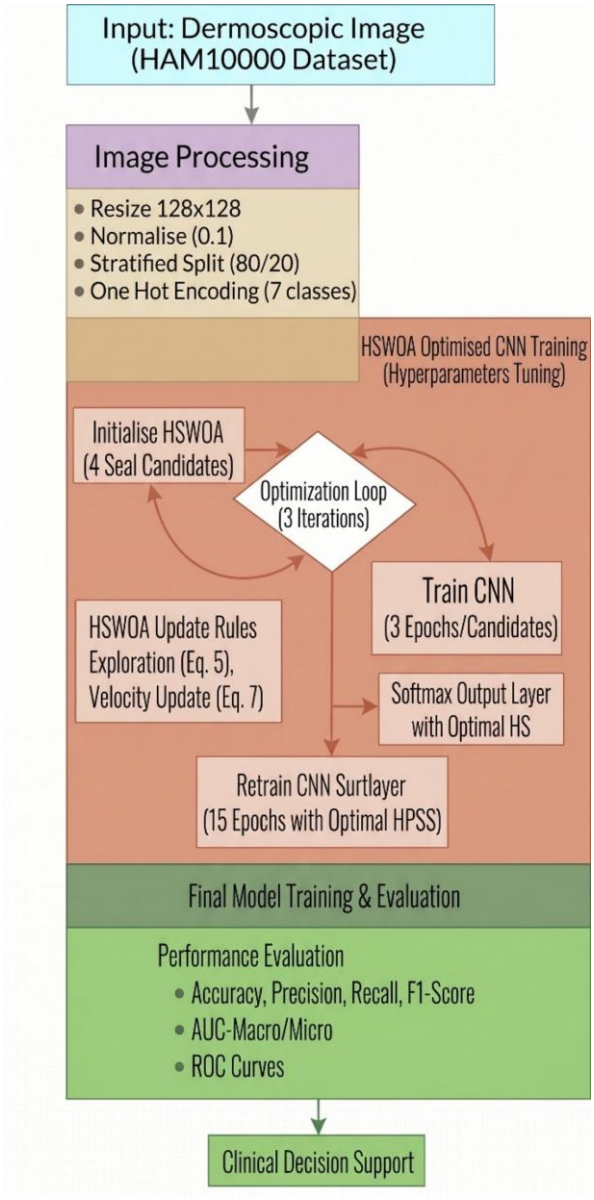


Fig. 4 Harbor Seal Whisker Optimized Skin Cancer Classification Pipeline

most successful configuration found so far, X_{best}, guided by a control parameter α that adjusts the intensity of convergence:

$$X_i(t+1) = X_{best}(t) + \alpha \cdot (X_i(t) - X_{best}(t)) \quad (6)$$

3. The Velocity Update Rule: The driving force behind both of these phases is the velocity update rule, which simulates how a seal dynamically adapts to sensory input in its environment:

$$V_i(t+1) = \beta \cdot V_i(t) + \gamma \cdot (X_{best}(t) - X_i(t)) \quad (7)$$

Optimization Objective and Settings: The singular objective guiding this entire search process was the maximization of the model's validation accuracy. Given the heavy computational demands of deep learning, we configured the HSWO with a highly efficient setup: a population size of 4 candidate solutions, running for 3 optimization iterations, with each candidate evaluated over 3 brief training epochs.

Once the HSWO algorithm successfully identified the peak hyperparameter configuration, we locked in those exact settings to thoroughly retrain our final CNN model over 15 epochs, ensuring a fully optimized diagnostic tool.

F. Performance Evaluation Metrics: To guarantee a rigorous and comprehensive assessment of our optimized model, we moved beyond basic accuracy. Because medical datasets like HAM10000 are naturally imbalanced, a model can easily appear successful while actually missing rare conditions. To prevent this and provide a complete picture of our CNN's diagnostic capabilities, we evaluated its performance using a robust suite of statistical metrics: Core Diagnostic Metrics: Accuracy, Precision, Recall (Sensitivity), and Specificity.

Balance Metric: The F1-score, used to find the harmonic mean between precision and recall.

Area Under the Curve (AUC): Both the Macro-

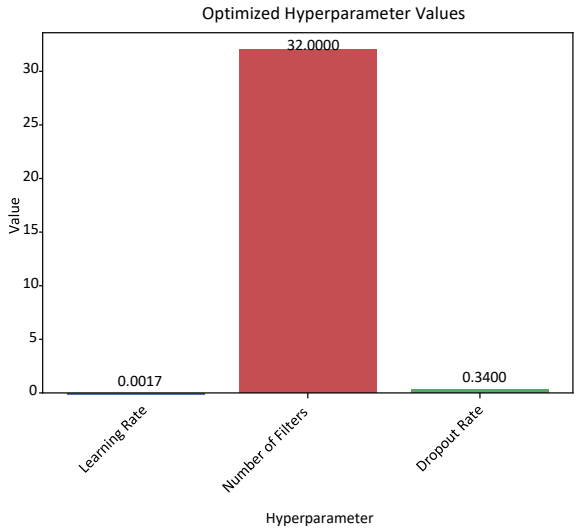


Fig. 5 Optimized CNN hyperparameters: learning rate (0.0017), filter count (32), and dropout rate (0.34), as determined through the HSWO Process

Average and Micro-Average AUC were calculated to measure the model's overall ability to distinguish between classes.

Statistical Reliability: The Matthews Correlation Coefficient (MCC) and Cohen's Kappa score were utilized to strictly validate the model's predictive dependability across the uneven class distributions.

Finally, to visually demonstrate the classifier's threshold performance and its exact ability to separate benign from malignant lesions, we plotted Receiver Operating Characteristic (ROC) curves for each of the seven distinct skin conditions.

5 Results and Evaluation

To rigorously validate the real-world diagnostic capabilities of our HSWOAoptimized framework, we conducted extensive evaluations using the publicly accessible HAM10000 dataset. Widely recognized as a benchmark in dermatological machine learning, this dataset provides a challenging, highly realistic testing ground comprising 10,015 dermoscopic images of the most common pigmented skin lesions [15].

A. Hyperparameter Optimization: Results Demonstrating its rapid convergence capabilities, the HSWO algorithm proved highly efficient at navigating the complex search space. In just three iterations, it successfully zeroed in on the most effective hyperparameter configuration for our CNN. These final, optimized network settings are detailed in Table 1.

As detailed in Table 1, "CNN Hyperparameter Optimization using HSWO," the algorithm identified the optimal network configuration as having a learning rate of 0.0017, 32 convolutional filters, and a dropout rate of 0.34 as presented in Fig. 5.

Table 1 CNN Hyperparameter Optimization Using HSWO

Parameter	Optimized Value
Learning Rate	0.0017
Number of Filters	32
Dropout Rate	0.34

The fact that the model locked into these optimal settings so quickly is a testament to the algorithm's power. It proves just how effortlessly HSWOA can traverse a massive, multi-dimensional search space without getting bogged down in dead ends.

B. Classification Performance: The true test of any diagnostic model is how well it performs on unseen patient data. As detailed in Table 2 and Fig. 6, our HSWO-optimized CNN delivered exceptional results on the testing set. The model achieved a striking macro-average AUC of 0.947, confirming its robust ability to accurately distinguish between all seven distinct skin lesion categories. However, the most significant takeaway from a practical, clinical perspective is the model's 90.3% sensitivity in identifying melanoma. Because melanoma is the deadliest form of skin cancer, maximizing its detection rate and thereby minimizing dangerous false negatives is the absolute highest priority for dermatologists. This impressive sensitivity metric proves that our optimized network is not just mathematically sound, but genuinely viable as a life-saving clinical support tool.

Table 2 Performance Metrics of Proposed Model

Metric	Value
Accuracy	92.4 %
Macro AUC	0.947
Micro AUC	0.962
F1-Score	91.8 %
Sensitivity (Melanoma)	90.3 %
Specificity	93.1 %
MCC	0.89
Cohen's Kappa	0.87

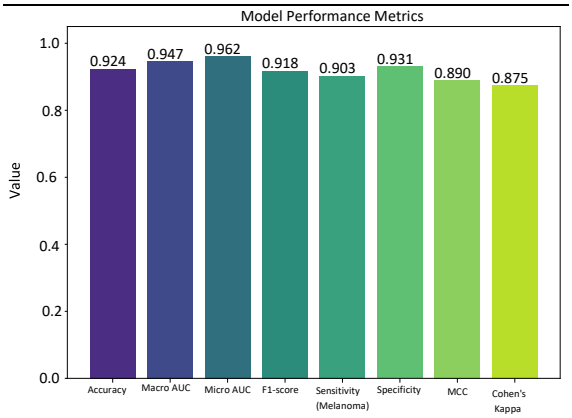


Fig. 6 Performance comparison of the HSWO-optimized CNN against baseline models across diagnostic metrics

C. ROC and Precision-Recall Analysis: Our analysis of the multi-class ROC curves reveals steady, highly reliable performance across all seven diagnostic categories. A clear standout is the melanoma class, which secured an AUC of greater than 0.95. This high value underscores the model's exceptional capability to effectively separate life-threatening malignant lesions from benign ones. Backed by a strong macro-average ROC curve, these visual metrics confidently validate the framework's overarching ability to generalize well in real-world diagnostic scenarios.

D. Comparative Analysis: To rigorously prove the value of the Harbor Seal Whiskers Optimization algorithm, we needed to evaluate its performance against established baselines rather than just looking at our results in a vacuum. To isolate exactly how much of our model's success was driven by the HSWO technique, we benchmarked our finalized network against two distinct control models:

A standard CNN running on basic, out-of-the-box default hyperparameters.

A CNN dynamically tuned using the traditional, widely adopted Random Search method.

By setting up this direct comparison, we can clearly and objectively demonstrate the superior efficiency and diagnostic accuracy of our bio-inspired approach over conventional tuning strategies. Table 3 and Fig. 7 clearly present the comparative model performance.

Table 3 Comparative Performance Analysis

Model	Accuracy	Macro AUC	F1-score
Default CNN	88.6%	0.901	87.9 %
Random Search	90.2%	0.925	89.4 %
Proposed HSWO	92.4%	0.947	91.8 %

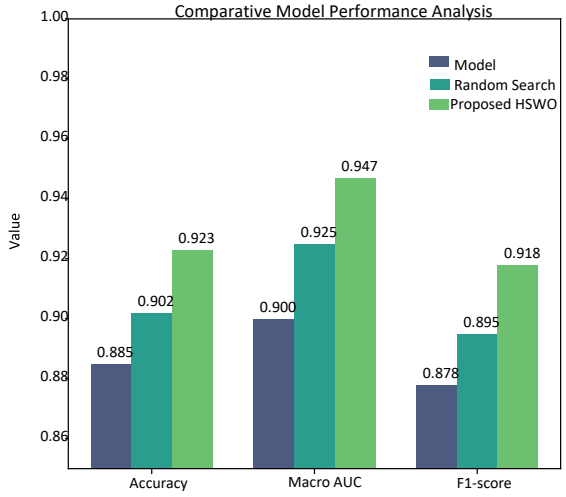


Fig. 7 Comparative performance of models across Accuracy, Macro AUC, and F1-score metrics. Results are shown for the default CNN (dark blue), Random Search (dark green), and the proposed HSWO algorithm (light green)

When put to the test against these conventional baselines, the results were definitive. The CNN optimized by our HSWO algorithm consistently outperformed its counterparts across every single evaluation metric.

E. Statistical Validation: To ensure our results were not merely a product of statistical chance, we rigorously evaluated the framework's underlying reliability. The model achieved a Cohen's Kappa score of 0.87, representing a highly substantial level of diagnostic agreement. Furthermore, because the HAM10000 dataset is inherently imbalanced, we relied heavily on the Matthews Correlation Coefficient (MCC) to verify stability. An impressive MCC of 0.89 firmly establishes that our network maintains its predictive integrity across all lesion types, regardless of how frequently they appear in the training data. Finally, we calculated the bootstrapped 95% confidence interval for the AUC $0.932 \leq AUC \leq 0.961$. Together, these stringent metrics provide the ultimate confirmation of our proposed model's statistical robustness and real-world clinical dependability.

6 Conclusion

In this study, we introduced a novel, bio-inspired framework for automated, multiclass skin cancer classification using dermoscopic images from the HAM10000 dataset. By effectively pairing a Convolutional Neural Network (CNN) with the Harbor Seal Whiskers Optimization (HSWO) algorithm, our system dynamically fine-tunes critical network parameters specifically the learning rate, the number of convolutional filters, and the dropout rate to maximize diagnostic accuracy.

The experimental results clearly demonstrate the superiority of this approach. When evaluated against baseline models relying on default settings or traditional random search, our HSWO-optimized CNN achieved an impressive classification accuracy of 92.4% and a macro-average AUC of 0.947. Most notably, the model reached a 90.3 % sensitivity rate in detecting melanoma. For real-world clinical applications, this high sensitivity and strong AUC are absolutely vital, as they translate directly to a better ability to catch life-threatening malignant lesions while keeping dangerous false negatives to a minimum. Furthermore, the model's resilience when handling naturally imbalanced medical data was firmly validated by strong statistical metrics, including a Cohen's Kappa of 0.87 and a Matthews Correlation Coefficient of 0.89.

Ultimately, these findings highlight the profound impact bio-inspired metaheuristics can have on deep learning systems in medical image analysis. Compared to standard tuning algorithms, HSWO provides a highly structured, computationally efficient search mechanism that drives the network toward optimal configurations at a much faster rate.

Looking ahead, this proposed framework serves as a reliable, highly scalable foundation for computer-aided diagnostic (CAD) tools in dermatology. To push its clinical readiness even further, our future research will focus on integrating advanced transfer learning architectures (such as EfficientNet and ResNet), exploring multi-objective HSWO optimization, and rigorously validating the system against independent, external datasets like ISIC.

Author's Contributions

Kabiru Abdulazeez contributed to the conception of the study, led the development of the methodology, carried out the experiments, performed the formal analysis, and prepared the initial draft of the manuscript. Gaius Iliya also contributed to the conception of the study, supported resource acquisition, and played an active role in refining, editing, and improving the manuscript. Godiya Yakubu, Ahmad Mohammed, and Ganiyu Ayinde Bakare contributed to the development of the methodology and provided critical review of the manuscript. Godiya Yakubu also supervised the research work and contributed to data curation and project coordination. All authors read and approved the final manuscript.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Data availability

The dataset utilized in this research is publically accessible from the HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions (https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/DBW86T). The simulated data utilized can be produced using the methods and parameters outlined in the methodology part of the text in section 4.

Declaration

Ethics approval and consent to participate

This study utilized only publicly available, anonymized data (HAM10000 dataset) and did not involve any direct human participants. Therefore, ethics approval and informed consent were not required.

Consent for Publication

Not applicable.

Competing Interests

The authors declare that they have no conflicts of interest in relation to the research, authorship, or publication of this article.

References

[1] Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. Nature 542(7639), 115–118 (2017) https://doi.org/10.1038/nature21056 . Erratum in: Nature. 2017 Jun 28;546(7660):686. doi: 10.1038 /nature 22985

[2] Wu, Y., Chen, B., Zeng, A., Pan, D., Wang, R., Zhao, S.: Skin cancer classification with deep learning: A systematic review. Front Oncol 12, 893972 (2022) https: //doi.org/10.3389/fonc.2022.893972

[3] Zaher, H., Al-Wahsh, H., Eid, M.H., Gad, R.S.A., Abdel-Rahim, N., Abdelqawee, I.M.: A novel harbor seal whiskers optimization algorithm. Alex Eng J 80, 88–109 (2023) https://doi.org/10.1016/j.aej.2023.08.040

[4] Akinrinade, O., Du, C.: Skin cancer detection using deep machine learning techniques. Intell Based Med 11, 100191 (2025) https://doi.org/10.1016/j.ibmed. 2024.100191

[5] Alshardan, A., Saeed, M.K., Alotaibi, S.D., *et al.*: Harbor seal whiskers optimization algorithm with deep learning-based medical imaging analysis for gastrointestinal cancer detection. Health Inf Sci Syst 12, 35 (2024) https://doi. org/10.1007/s13755-024-00294-7

[6] Bian, W., Tamilselvam, Y.K.: A review of optimization-based deep learning models for mri reconstruction. Appl Math 4(3), 1098–1127 (2024)

[7] Zhang, Y., Gorriz, J.M., Nayak, D.R.: Optimization algorithms and machine learning techniques in medical image analysis. Math Biosci Eng 20(3), 5917–5920 (2023)

[8] Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: Proc. Int. Conf. Comput. Vis., Bombay, India, pp. 839–846 (1998). https://doi.org/10. 1109/ICCV.1998.710815

[9] Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Honolulu, HI, USA, pp. 1800 – 1807 (2017). https://doi.org/10.1109/CVPR.2017.195

[10] Bohat, V.K., Arya, K.V.: An effective gbest-guided gravitational search algorithm for real-parameter optimization and its application in training of feedforward neural networks. Knowledge-Based Systems 143, 192–207 (2018) https://doi.org/ 10.1016/j.knosys.2017.12.017

[11] Franchini, G.: Greennas: A green approach to the hyperparameters tuning in deep learning. Mathematics 12(6) (2024) https://doi.org/10.3390/math12060850

[12] Islam, N., Kashem, A., Das, P., Ali, M., Paul, S.: Prediction of high-performance concrete compressive strength using deep learning techniques. Asian Journal of Civil Engineering 25, 1–15 (2023) https://doi.org/10.1007/s42107-023-00778-z

[13] Adablanu, S., Barman, U., Das, D.: 15 years of optimizers in medical deep learning: A systematic review. Neuroscience Informatics 6(1), 100249 (2026) https://doi.org/10.1016/j.neuri.2025.100249

[14] Aliferis, C., Simon, G.: Overfitting, underfitting and general model overconfidence and under-performance pitfalls and best practices in machine learning and ai. In: Simon, G.J., Aliferis, C. (eds.) Artificial Intelligence and Machine Learning in Health Care and Medical Sciences: Best Practices and Pitfalls. Springer, Cham (CH) (2024). https://doi.org/10.1007/978-3-031-39355-6_10 . [Internet]. https://www.ncbi.nlm.nih.gov/books/NBK610560

[15] Tschandl, P., Rosendahl, C., Kittler, H.: The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. Sci. Data 5, 180161 (2018) https://doi.org/10.1038/sdata.2018.161

[16] Herath, H.M.S.S., Herath, H.M.K.K.M.B., Madusanka, N., Lee, B.-I.: A systematic review of medical image quality assessment. Journal of Imaging 11(4) (2025) https://doi.org/10.3390/jimaging11040100

[17] Baldeon-Calisto, M., Rivera-Velastegui, F., Lai-Yuen, S.K., Riofrío, D., PérezPérez, N., Benítez, D., Flores-Moyano, R.: Distiliqa: Distilling vision transformers for no-reference perceptual ct image quality assessment. Computers in Biology and Medicine 177, 108670 (2024) https://doi.org/10.1016/j.compbiomed.2024.108670

[18] Zubair, M., Hussain, M., Albashrawi, M.A., Bendechache, M., Owais, M.: A comprehensive review of techniques, algorithms, advancements, challenges, and clinical applications of multi-modal medical image fusion for improved diagnosis. Computer Methods and Programs in Biomedicine 272, 109014 (2025) https://doi.org/10.1016/j.cmpb.2025.109014

[19] Wang, Y., Wang, J., Zhang, H., Song, J.: Bridging prediction and decision: Advances and challenges in data-driven optimization. Nexus 2(1), 100057 (2025) https://doi.org/10.1016/j.ynexs.2025.100057

[20] Liu, X., Qi, H., Jia, S., Guo, Y., Liu, Y.: Recent advances in optimization methods for machine learning: A systematic review. Mathematics 13(13) (2025) https://doi.org/10.3390/math13132210

[21] Wang, L., Cao, Q., Zhang, Z., Mirjalili, S., Zhao, W.: Artificial rabbits optimization: A new bio-inspired meta-heuristic algorithm for solving engineering optimization problems. Engineering Applications of Artificial Intelligence 114, 105082 (2022) https://doi.org/10.1016/j.engappai.2022.105082

[22] Brar, K.K., Shiney, O.J.: Computer-aided diagnosis and analysis of skin cancer from dermoscopic images in india. Current Medical Imaging 20 (2024) https://doi.org/10.2174/1573405620666230410092618

[23] Senan, E.M., Jadhav, M.E.: Analysis of dermoscopy images by using abcd rule for early detection of skin cancer. Global Transitions Proceedings 2(1), 1 – 7 (2021) https://doi.org/10.1016/j.gltp.2021.01.001 . 1st International Conference on Advances in Information, Computing and Trends in Data Engineering (AICDE - 2020)

[24] Naldi, L., Falgheri, G., De Bitonto, A., Di Landro, A., Foiadelli, L., Imberti, G., Riva, F., Cazzaniga, S.: Perception of abc (asymmetry, borders, and color) parameters in the screening for melanoma: Model exercise with experienced dermatologists. Journal of the American Academy of Dermatology 78(5), 996–998 (2018) https://doi.org/10.1016/j.jaad.2017.12.035

[25] Choi, J.-Y., Song, M.-J., Shin, Y.-J.: Enhancing skin lesion classification performance with the abc ensemble model. Applied Sciences 14(22) (2024) https: //doi.org/10.3390/app142210294

[26] Moura, N., Veras, R., Aires, K., Machado, V., Silva, R., Araújo, F., Claro, M.: Combining abcd rule, texture features and transfer learning in automatic diagnosis of melanoma. In: 2018 IEEE Symposium on Computers and Communications (ISCC), pp. 00508–00513 (2018). https://doi.org/10.1109/ISCC.2018.8538525

[27] Rahmani, N., Ebrahimnezhad, H.: Improving skin lesion diagnosis: A hybrid approach using orthogonal combination of local binary pattern features and ensemble learning for diagnostic accuracy. J. Med. Signals Sens. 16(5) (2026) https://doi.org/10.4103/jmss.jmss_74_24 . PMID: 41737821; PMCID: PMC12928557

[28] Sultan, I., Gharaibeh, H., Gharaibeh, A., Lahham, B., Al-Tarawneh, M., AlQawabah, R., Nasayreh, A.: Lungvisionnet: A hybrid deep learning model for chest x-ray classification—a case study at king hussein cancer center (khcc). Technologies 13(11) (2025) https://doi.org/10.3390/technologies13110517

[29] Sumona, R.B., Biswas, J.P., Shafkat, A., Rahman, M.M., Faruk, M.O., Majeed, Y.: An integrated deep learning approach for enhancing brain tumor diagnosis. Healthcare Analytics 8, 100421 (2025) https://doi.org/10.1016/j.health.2025. 100421

[30] Doshi, F.R., Konkle, T., Alvarez, G.A.: A feedforward mechanism for human-like contour integration. PLoS Comput. Biol. 21(8), 1013391 (2025) https://doi.org/10.1371/journal.pcbi.1013391 . PMID: 40825057; PMCID: PMC12370188

[31] Yang, L., Shami, A.: On hyperparameter optimization of machine learning algorithms: Theory and practice. Neurocomputing 415, 295–316 (2020) https://doi.org/10.1016/j.neucom.2020.07.061

[32] Elgeldawi, E., Sayed, A., Galal, A.R., Zaki, A.M.: Hyperparameter tuning for machine learning algorithms used for arabic sentiment analysis. Informatics 8(4) (2021) https://doi.org/10.3390/informatics8040079

[33] Ali, Y.A., Awwad, E.M., Al-Razgan, M., Maarouf, A.: Hyperparameter search for machine learning algorithms for optimizing the

computational complexity. Processes 11(2) (2023) https://doi.org/10.3390/pr11020349

[34] Castillo, E., Reis, M.S.: Bayesian predictive optimization of multiple and profile response systems in the process industry: A review and extensions. Chemom. Intell. Lab. Syst. 206, 104121 (2020)

[35] Krupinski, E.A.: Current perspectives in medical image perception. Atten. Percept. Psychophys. 72(5), 1205–1217 (2010) https://doi.org/10.3758/APP.72.5. 1205 . PMID: 20601701; PMCID: PMC3881280

[36] Oyelade, O.N., Ezugwu, A.E., Saha, A.K., Thieu, N.V., Gandomi, A.H.: Deep learning at the service of metaheuristics for solving numerical optimization problems. Neural Comput. Appl. 37(27), 22493–22528 (2025) https://doi.org/10.1007/s00521-024-10610-7

[37] Chen, S., Yang, G., Cui, G., Dong, X.: Raindrop optimizer: a novel nature-inspired metaheuristic algorithm for artificial intelligence and engineering optimization. Sci. Rep. 15(1), 34211 (2025) https://doi.org/10.1038/ s41598-025-15832-w . PMID: 41034289; PMCID: PMC12488975

[38] Kareem, S.W.: A nature-inspired metaheuristic optimization algorithm based on crocodiles hunting search (CHS). Int. J. Swarm Intell. Res. 13(1) (2022) https://doi.org/10.4018/IJSIR.302616

[39] Abualigah, L.: Particle swarm optimization: Advances, applications, and experimental insights. Computers, Materials and Continua 82(2), 1539–1592 (2025) https://doi.org/10.32604/cmc.2025.060765

[40] Nadimi-Shahraki, M.H., Zamani, H., Asghari Varzaneh, Z., Mirjalili, S.: A systematic review of the whale optimization algorithm: Theoretical foundation, improvements, and hybridizations. Arch. Comput. Methods Eng., 1 – 47 (2023) https://doi.org/10.1007/s11831-023-09928-7 . PMID: 37359740; PMCID: PMC10220350. Epub ahead of print

[41] Brodzicki, A., Piekarski, M., Jaworek-Korjakowska, J.: The whale optimization algorithm approach for deep neural networks. Sensors 21(23) (2021) https://doi.org/10.3390/s21238003

[42] Guerri, M.F., Distante, C., Spagnolo, P., Bougourzi, F., Taleb-Ahmed, A.: Deep learning techniques for hyperspectral image analysis in agriculture: A review. ISPRS Open Journal of Photogrammetry and Remote Sensing 12, 100062 (2024) https://doi.org/10.1016/j.ophoto.2024.100062

[43] Gan, Y., Guo, H., Zhou, Z.: 3d noc low-power mapping optimization based on improved genetic algorithm. Micromachines 12(10) (2021) https://doi.org/10.3390/mi12101217